

SGI keeps storage down to earth

LABS ON-SITE: DMF SYSTEM HAS SCALED UP TO HANDLE NASA UNIT'S ASTRONOMICAL DATA LOADS

By Henry Baltazar

WITH HIS USERS GENERATING DATA AT INCREDIBLE RATES, ALAN POWERS, high-end-computing lead for NASA's Advanced Supercomputing Division, had to find a way to efficiently and inexpensively store data while keeping it accessible to users.

About six years ago, Powers tested Silicon Graphics Inc.'s DMF (Data Migration Facility) solution, which blends tape- and hard-drive-based storage using automated data migration policies. Impressed enough during the tests, he decided to replace the Advanced Supercomputing Division's homegrown HSM

(hierarchical storage management) system with the DMF.

eWEEK Labs recently got a look at the system, which Powers said has capably handled everything the division's users have thrown at it over the years.

SGI's DMF, a data life-cycle management solution, is designed for high-performance computing environments. Using DMF, IT managers can grow storage capacity by augmenting primary storage in the form of tapes and inexpensive disks such as Serial ATA RAID setups (see chart).

By going with a DMF system instead of just adding more hard drive RAIDs, Powers estimates that he has saved five to 10 times the cost of purchasing more RAIDs.

"SGI's DMF came into existence about six years ago," Powers said. "Within six months after SGI announced it, we had it up, and while we were evaluating it, we made a decision that we were going to replace our custom HSM solution ... with DMF." Since then, the division has added more servers to accommodate growing data loads, he said.

Over the years, Powers' DMF implementation has scaled up nicely. "In October of 1999, we were archiving 160GB

of data a day [on the SGI system]. Now our average is 1,500GB a day," he said.

Although Powers' site has workloads and operations that are a bit different from what we see in most standard IT shops, his use of data life-cycle management technology is applicable to non-research-oriented data centers.

The scalability Powers has seen in his system and the relative cost savings (over primary disk) are attractive goals that many IT managers continue to seek.

The Advanced Supercomputing Division, based at Moffitt Field, Calif., devel-


ops distributed heterogeneous computing capabilities to enable NASA projects and missions. The division is also charged with researching, developing and delivering high-end computing services and technologies, such as applications and algorithms, tools, and system software and hardware to NASA and others.

Academic users include the University of Glasgow, Massachusetts Institute of Technology, Stanford University and the University of Tennessee. Other users include the Department of Defense's Major Shared Resource Centers, The Boeing Co. and Lockheed Martin Corp.

Users routinely generate from 1 terabyte to 3 terabytes of data per day. To manage the data load, DMF automatically and seamlessly moves stale (not recently accessed) data from expensive Fibre Channel storage systems to a less expensive near-line tape library.

Powers said that one of his biggest challenges was fine-tuning the DMF system to make sure it would not hinder his users. However, because DMF leaves a

DMF puts data in its place—close to users



SGI's DMF solution allows IT managers to create a tiered storage infrastructure where older, less frequently accessed files are moved from expensive Fibre Channel RAID storage to less costly tape media. Users write data to a Linux or Irix server running DMF; the data is written to the primary storage. DMF's customizable migration policies determine when and where to migrate data. For example, an IT manager can configure DMF to move larger files to high-capacity, high-speed single-reel tapes while smaller files are moved to quick-searching single-reel tapes. For the user, this storage migration is nearly seamless. When a user requests a file that is in an archive, the DMF server finds and restores the file on the fly.

file pointer in place of the original file, the data migration process is usually undetectable to users, he said.

In addition, DMF's customization capabilities enabled Powers and his staff to easily optimize their data archive project. "One of the nice things about DMF is that you can customize it to basically your heart's content," said Powers. "DMF has generic and easy customizations that allowed us to segment users and groups and have separate policies for them."

Using the customization capabilities, Powers implemented data migration rules that tailored DMF to meet his users' needs. For example, Powers and his staff implemented a rule that directs DMF to store files that are less than 1MB in size on fast primary storage because doing so lets users quickly access source files.

Most users wait a little less than a minute per file, so the latency of the DMF systems has not harmed users' productivity. For retrieval of fairly large files (greater than 1GB), users might need to wait 5 to 10 minutes to get their files restored from the tape silos.

For the most part, Powers' users have been happy with the DMF solution. Powers said he has gotten a couple of complaints from users trying to access extremely large numbers of files or when there are network or silo problems.

Powers' DMF implementation uses a couple of different tape drives in his Storage Technology Corp. StorageTek 9310 tape silos.

For smaller files, the DMF archives data onto dual-reel tapes (StorageTek T9840A and T9840B tape drives), which have relatively low capacity (20GB per tape) but

speedy load (4 seconds) and search (8 seconds) times. The T9840A transfers data at 10MB per second; the T9840B transfers data at 19MB per second. These fast load and search times ensure that users don't have to wait long to access small files.

For larger files, the DMF uses StorageTek T9940B tape drives, which have higher capacities (200GB per tape) and higher data transfer rates (30MB per second). The drawback to these tapes is that they have long load (18 seconds on average) and search (41 seconds on average) times.

Using these two tape technologies, Powers was able to create a flexible tape archive that works well with small and large files. [e](#)

Senior Analyst Henry Baltazar can be reached at henry_baltazar@ziffdavis.com.

Reprinted from eWEEK, March 8, 2004 with permission from Ziff Davis Media Inc.
©2004 Ziff Davis Publishing Holdings Inc. All rights reserved.

The logo for StorageTek (SGI) consists of the lowercase letters 's', 'g', and 'i' in a bold, sans-serif font. The 's' and 'g' are connected at the bottom, and the 'i' has a dot above it. A registered trademark symbol (®) is located to the upper right of the 'i'.

www.sgi.com/storage